

Quadratic Optimization: Variant

Last time: we discussed finding the extremal (min & max) values of a quadratic form

$$q(x) = \sum_{i,j} a_{ij} x_i x_j$$

subject to the constraint $1 = \|x\|^2 = x_1^2 + \dots + x_n^2$.

Procedure: $q(x) = x^T S x$ for S symmetric

orthogonally diagonalize: $S = Q D Q^T$ $D = \begin{pmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{pmatrix}$

change variables: $x = Q y$

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$$

$$\Rightarrow q(x) = \lambda_1 y_1^2 + \dots + \lambda_n y_n^2$$

Answer:

maximum = λ_1 , achieved at any unit λ_1 -eigenvector

minimum = λ_n , achieved at any unit λ_n -eigenvector

Here's an (almost) equivalent variant of this problem that you can draw.

Quadratic Optimization Problem, Variant:

Given a quadratic form $q(x)$, find the minimum & maximum values of $\|x\|^2$ subject to $q(x) = 1$.

So we switched the **function** we're extremizing ($\|x\|^2$) and the **constraint** ($q(x)=1$).

In general the min/max may not exist.

- $q(x_1, x_2) = -x_1^2 - 2x_2^2$:

there is **no** x such that $q(x)=1$!


- $q(x_1, x_2) = x_1^2 - x_2^2$:

there is **no maximum** $\|x\|^2$

subject to $q(x)=1$:

$$q(C, \sqrt{C^2-1}) = 1$$

for any (huge) C .



Problem: $q(x)$ may be 0 or negative!

Def: A quadratic form is **positive-definite** if $q(x) > 0$ for all $x \neq 0$.

NB: If $q(x) = x^T S x$ then q is positive-definite $\iff S$ is positive-definite: this is the **positive-energy** criterion.

In this case, the problem is equivalent to the previous one, as follows:

Recall: $q(cx) = c^2 q(x)$

Fact: If q is positive-definite then

u maximizes $q(u)$
subject to $\|u\|=1$
with maximum
value λ_1

\Leftrightarrow

$x = \frac{1}{\sqrt{\lambda_1}} u$ minimizes
 $\|x\|^2$ subject to
 $q(x)=1$ with minimum
value $1/\lambda_1$.

and

u minimizes $q(u)$
subject to $\|u\|=1$
with minimum
value λ_n

\Leftrightarrow

$x = \frac{1}{\sqrt{\lambda_n}} u$ maximizes
 $\|x\|^2$ subject to
 $q(x)=1$ with maximum
value $1/\lambda_n$

Why? if $q(u) = \lambda > 0$ and $x = \frac{1}{\sqrt{\lambda}} u$ then
 $q(x) = q\left(\frac{1}{\sqrt{\lambda}} u\right) = \frac{1}{\lambda} q(u) = \frac{1}{\lambda} \cdot \lambda = 1$.

If λ is maximized then $\|x\|^2 = \frac{1}{\lambda}$ is minimized
and vice-versa.

So we know exactly how to solve this QO problem variant: do the same procedure as in the original QO problem, and take reciprocals.

Eg: Extremize $\|x\|^2$ subject to

$$q(x_1, x_2) = \frac{5}{2}x_1^2 + \frac{5}{2}x_2^2 - x_1x_2 = 1$$

Last time: $q(x) = x^T S x$

$$S = \frac{1}{2} \begin{pmatrix} 5 & -1 \\ -1 & 5 \end{pmatrix} = Q D Q^T \quad Q = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix} \quad D = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}$$

- q is maximized (subject to $\|u\|=1$)

at $u_1 = \pm \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix}$ with maximum value 3.

$$q(u_1) = 3 \quad x_1 = \frac{u_1}{\sqrt{3}} \Rightarrow q(x_1) = 1 \quad \|x_1\|^2 = \frac{1}{3}$$

The minimum value of $\|x\|^2$ subject to $q(x)=1$ is $1/3$. It is achieved at $x_1 = \frac{1}{\sqrt{3}} u_1$.

- q is minimized at $u_2 = \pm \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ with minimum value 2.

$$q(u_2) = 2 \quad x_2 = \frac{1}{\sqrt{2}} u_2 \Rightarrow q(x_2) = 1 \quad \|x_2\|^2 = \frac{1}{2}$$

The maximum value of $\|x\|^2$ subject to $q(x)=1$ is $1/2$. It is achieved at $x_2 = \frac{1}{\sqrt{2}} u_2$.


→ note $\frac{1}{2} > \frac{1}{3}$ ✓

Geometric Interpretation

Recall: An equation of the form

$$\lambda_1 x_1^2 + \lambda_2 x_2^2 = 1$$


$(\lambda_1, \lambda_2 > 0)$ defines an ellipse.



(This is a circle horizontally stretched by $1/\sqrt{\lambda_1}$, & vertically stretched by $1/\sqrt{\lambda_2}$)

If $q(x_1, x_2) = \lambda_1 x_1^2 + \lambda_2 x_2^2$ is diagonal & positive-definite then $q(x_1, x_2) = 1$ defines the ellipse above, and

extremizing $\|x\|^2 = 1$ subject to $q(x) = 1$ amounts to finding the shortest ($\pm x$) & longest ($\pm y$) vectors on the ellipse.



$$\|x\|^2 = 1/\lambda_1$$
$$\|y\|^2 = 1/\lambda_2$$

In general, $q(x) = \lambda_1 x_1^2 + \dots + \lambda_n x_n^2$ (all $\lambda_i > 0$) defines an ellipsoid ("egg"); extremizing $\|x\|^2$ subject to $q(x) = 1$ means finding the shortest & longest vectors.

Non-diagonal case:

$q(x) = x^T S x$ for S positive-definite.


Let $\lambda_1 \geq \lambda_2 > 0$ be the eigenvalues, u_1, u_2 orthonormal eigenvectors.

Change variables: $x = Qy$ $Q = \begin{pmatrix} u_1 & u_2 \end{pmatrix}$

$$\lambda_1 y_1^2 + \lambda_2 y_2^2 = 1$$




$$q(x) = 1$$



(y_1, y_2) -plane

multiply
by Q^T

$$\begin{aligned} u_1 &= Q e_1 \\ u_2 &= Q e_2 \end{aligned}$$



(x_1, x_2) -plane

Upshot: $q(x) = 1$ defines a (rotated) ellipse.

The minor axis is in the u_1 -direction.

→ The shortest vectors are $\pm \frac{1}{\sqrt{\lambda_1}} u_1$

The major axis is in the u_2 -direction.

→ The longest vectors are $\pm \frac{1}{\sqrt{\lambda_2}} u_2$.

So we've drawn a picture of quadratic optimization problem (variant).


Everything works in higher dimensions; just get rotated ellipsoids.

Eg: $q(x_1, x_2) = \frac{5}{2}x_1^2 + \frac{5}{2}x_2^2 - x_1x_2 = x^T S x$

$$S = \frac{1}{2} \begin{pmatrix} 5 & -1 \\ -1 & 5 \end{pmatrix} = Q D Q^T \quad Q = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix} \quad D = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}$$

$$x = Q y \rightarrow q = 3y_1^2 + 2y_2^2$$

$$3y_1^2 + 2y_2^2 = 1$$




(y_1, y_2) -plane

multiply
by Q

$$\begin{aligned} u_1 &= Q e_1 \\ u_2 &= Q e_2 \end{aligned}$$

$$q(x) = 1$$



$$\begin{aligned} u_1 &= \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix} \\ u_2 &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{aligned}$$

(x_1, x_2) -plane

The orthogonal diagonalization procedure took the ellipse

$$q(x_1, x_2) = \frac{5}{2}x_1^2 + \frac{5}{2}x_2^2 - x_1x_2$$

and found its major & minor axes & radii.

Additional Constraints

These come up naturally in practice (see the spectral graph theory problem on HW13) and in the PCA.

"Second-largest" value:

Suppose $q(x)$ is maximized (subject to $\|x\|=1$) at u_1 .
What is the maximum value of $q(x)$ subject to
 $\|x\|=1$ and $x \perp u_1$?

This rules out the maximum value \rightarrow get "second-largest" value.

How to solve this?

- Write $q(x) = x^T S x$
- Orthogonally diagonalize $S = Q D Q^T$

Suppose u_1 is the first column of Q (1^{st} λ_1 -eigenvector)

- Set $x = Q y$

$$q = \lambda_1 y_1^2 + \dots + \lambda_n y_n^2 \quad \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$$

Answer: The maximum value of $q(x)$ subject to $\|x\|=1$ & $x \perp u_1$ is λ_2 . It is achieved at any unit λ_2 -eigenvector u_2 that is $\perp u_1$.

NB: If $\lambda_1 > \lambda_2$ then $u_2 \perp u_1$ automatically.

Why?

- If $q = \lambda_1 y_1^2 + \dots + \lambda_n y_n^2$ is diagonal then $u_1 = e_1 = (1, 0, \dots)$ so $x \perp u_1$ means $y_1 = 0$ \leadsto extremizing $\lambda_2 y_2^2 + \lambda_3 y_3^2 + \dots + \lambda_n y_n^2$.
- Otherwise, change variables $x = Qy$.

Q is orthogonal, so

$$y \cdot e_1 = 0 \iff 0 = (Qy) \cdot (Qe_1) = x \cdot u_1$$

$$\|y\| = 1 \iff 1 = \|Qy\| = \|x\|$$

(relate constraints on x & y)

Eg: Find the largest and second-largest values of $q(x) = 2x_1^2 + 2x_2^2 + 5x_3^2 + 2x_1x_2 - 8x_1x_3 + 8x_2x_3$ subject to $x_1^2 + x_2^2 + x_3^2 = 1$.

- $q = x^T S x$ for $S = \begin{pmatrix} 2 & 1 & -4 \\ 1 & 2 & 4 \\ -4 & 4 & 5 \end{pmatrix}$

- $S = Q D Q^T$ for

$$Q = \begin{pmatrix} -1/\sqrt{6} & 1/\sqrt{2} & 1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{2} & -1/\sqrt{3} \\ 2/\sqrt{6} & 0 & 1/\sqrt{3} \end{pmatrix} \quad D = \begin{pmatrix} 9 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & -3 \end{pmatrix}$$

Largest value is $q(x)=9$ at $x = \pm \frac{1}{\sqrt{6}} \begin{pmatrix} -1 \\ 1 \\ 2 \end{pmatrix} = \pm u_1$

Second-largest value:

The maximum value of $q(x)$ subject to

$\|x\|=1$ & $x \perp u_1$ is

$q(x)=3$ achieved at $x = \pm \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$

This also works for minimizing.

Second-smallest value:

Suppose $q(x)$ is minimized (subject to $\|x\|=1$) at u_n .

What is the minimum value of $q(x)$ subject to

$\|x\|=1$ and $x \perp u_n$?

Answer: The minimum value of $q(x)$ subject to $\|x\|=1$ & $x \perp u_n$ is λ_{n-1} . It is achieved at any unit λ_{n-1} -eigenvector u_{n-1} that is $\perp u_n$.

(automatic if $\lambda_{n-1} > \lambda_n$)

You can keep going:

Third-largest value:

Suppose $q(x)$ is maximized (subject to $\|x\|=1$) at u_1
and $q(x)$ is maximized (subject to $\|x\|=1$ and $x \perp u_1$)
at u_2 .

What is the maximum value of $q(x)$ subject to
 $\|x\|=1$ and $x \perp u_1$ and $x \perp u_2$?

NB: This "rules out" the largest & second-largest values.

Answer: The maximum value of $q(x)$ subject to

$\|x\|=1$ & $x \perp u_1$ & $x \perp u_2$ is λ_3 . It is achieved at
any unit λ_3 -eigenvector u_3 that is $\perp u_1$ and u_2 .

(automatic if $\lambda_2 > \lambda_3$)

This also works for the variant problem, except you
have to take reciprocals.

Et cetera...

Quadratic Optimization for $S=A^T A$

This is what we'll use for PCA.

Let $S=A^T A$ and $q(x)=x^T S x$. Then

$$\begin{aligned} q(x) &= x^T S x = x^T (A^T A) x = (x^T A^T) (A x) \\ &= (A x)^T (A x) = (A x) \cdot (A x) = \|A x\|^2 \end{aligned}$$

$$S=A^T A \Rightarrow x^T S x = \|A x\|^2$$

In this case, extremizing $q(x)$ subject to $\|x\|=1$ means extremizing $\|A x\|^2$ subject to $\|x\|=1$.

Procedure: to extremize $\|A x\|^2$ subject to $\|x\|=1$:

Orthogonally diagonalize $S=A^T A$

\hookrightarrow orthonormal eigenbasis $\{u_1, \dots, u_n\}$,
eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$.

- The largest value is λ_1 , achieved at any unit λ_1 -eigenvector u_1 .
- The smallest value is λ_n , achieved at any unit λ_n -eigenvector u_n .
- The second-largest value is λ_2 , achieved at any unit λ_2 -eigenvector $u_2 \perp u_1$ etc.

NB: these are eigenvectors/eigenvalues of $S = A^T A$, not of A (which need not be square).

Def: The **matrix norm** of a matrix A is

$\|A\|$ = the maximum value of $\|Ax\|$ subject to $\|x\| = 1$.

So $\|A\| = \sqrt{\lambda_1}$ λ_1 = largest eigenvalue of $A^T A$.

Eg: Compute $\|A\|$ for $A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix}$.

$$A^T A = \begin{pmatrix} 3 & 2 \\ 2 & 3 \end{pmatrix} \quad p(\lambda) = \lambda^2 - 6\lambda + 5 = (\lambda - 5)(\lambda - 1)$$

The largest eigenvalue is $\lambda = 5$, so $\|A\| = \sqrt{5}$.

$$\text{Eigenvector: } \begin{pmatrix} -b \\ a - \lambda \end{pmatrix} = \begin{pmatrix} -2 \\ -2 \end{pmatrix}$$

$$\text{Unit eigenvector: } u_1 = \frac{1}{\sqrt{2^2 + 2^2}} \begin{pmatrix} -2 \\ -2 \end{pmatrix} = \frac{1}{2\sqrt{2}} \begin{pmatrix} -2 \\ -2 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ -1 \end{pmatrix}$$

$$\text{Check: } Au_1 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ -1 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ -2 \\ -1 \end{pmatrix}$$

$$\text{has length } \frac{1}{\sqrt{2}} \cdot \sqrt{1^2 + 2^2 + 2^2 + 1^2} = \frac{\sqrt{10}}{\sqrt{2}} = \sqrt{5} \quad \checkmark$$